

Sparse Estimation with Generalized Beta Mixture and the Horseshoe Prior

Zahra Sabetsarvestani^{*}, Hamidreza Amindavar[†]

Abstract

In this paper, the use of the Generalized Beta Mixture (GBM) and Horseshoe distributions as priors in the Bayesian Compressive Sensing framework is proposed. The distributions are considered in a two-layer hierarchical model, making the corresponding inference problem amenable to Expectation Maximization (EM). We present an explicit, algebraic EM-update rule for the models, yielding two fast and experimentally validated algorithms for signal recovery. Experimental results show that our algorithms outperform state-of-the-art methods on a wide range of sparsity levels and amplitudes in terms of reconstruction accuracy, convergence rate and sparsity. The largest improvement can be observed for sparse signals with high amplitudes.

1 Introduction

Compressive Sensing is a new signal processing framework for efficiently acquiring and reconstructing a signal that have a sparse representation in a fixed linear basis. Consider the following acquisition system:

$$\mathbf{y} = \Phi \mathbf{w} + \mathbf{n} \quad (1)$$

where \mathbf{y} is the $M \times 1$ measurement vector, Φ is the $M \times N$ measurement matrix where $M < N$ and \mathbf{n} represents the measurement noise. In this expression \mathbf{w} is the unknown signal which is considered to be sparse. By exploiting the sparsity constraint on \mathbf{w} the inversion of (1) which is an ill-posed problem is the solution of the following optimization problem:

$$\hat{\mathbf{w}} = \arg \min_{\mathbf{w}} \|\mathbf{y} - \Phi \mathbf{w}\|_2 + \tau \|\mathbf{w}\|_0 \quad (2)$$

where $\|\mathbf{w}\|_0$ is equal to the number of non-zero entries of \mathbf{w} . This optimization problem is NP-Hard and can not be solved in polynomial time, but it has been shown in [1] that under some certain conditions on the measurement matrix Φ , (2) can be replaced by the following convex optimization problem.

$$\hat{\mathbf{w}} = \arg \min_{\mathbf{w}} \|\mathbf{y} - \Phi \mathbf{w}\|_2 + \tau \|\mathbf{w}\|_1 \quad (3)$$

Another class of CS recovery methods uses a greedy algorithm to estimate the sparse signal. In the greedy algorithms the support of the sparse signal is estimated sequentially and a local optimization in each step estimates the signal coefficients [2].

The main focus of this paper is on the Bayesian approaches in compressive sensing, known as the Bayesian Compressive Sensing (BCS). BCS which is based on the *Relevance Vector Machine* (RVM) or *Sparse Bayesian Learning* (SBL) theory [4, 5] was introduced in [3] for the first time. In BCS, all of the problem components are modeled in the Bayesian framework and a proper inference method is utilized to reconstruct the sparse signal. Developing a sparsity-inducing prior for modeling the sparse signal efficiently is a main research topic in BCS. The priors should be concentrated on zero to model the sparsity and at the

^{*}Amirkabir University of Technology. zahra.sabet@aut.ac.ir

[†]Amirkabir University of Technology. hamidami@aut.ac.ir

same time, have heavy tail to avoid over-shrinkage of larger signal coefficients. The recent work in this area was led to the introduction of new priors [11, 6]. In [6] a spike and slab prior is considered on the cluster sparse signal coefficients and a Markov Chain Monte Carlo (MCMC) sampling method is utilized for the inference. Since in MCMC a large set of samples are required for the posterior distribution estimation, this method suffers from high computational complexity and low convergence rate in consequence. Another class of sparsity-inducing priors which is widely used in literature [7, 3] exploit an equivalent conjugate hierarchy model. In the hierarchical model a complex prior is decomposed to two or more simple distributions which facilitate the inference procedure. In RVM [4] and in its adaption to BCS [3] a student's t-distribution is assigned as a two layer hierarchical prior on the sparse coefficients. In the later work [7] a hierarchical model of Laplace distribution is used. Laplace and student's t-distribution is concentrated near origin which shrink small coefficients toward zero, but due to having light tails, they over-shrink coefficients not close to zero and fail in modeling signals with large amplitude. large-amplitude coefficients arises in applications where signal shows impulsive behaviour [13].

In this paper we suggest to use *Generalized Beta Mixture* (GBM) and *Horseshoe* distributions as priors in the BCS framework. The GBM has two free parameters which promote its capability in modeling large-amplitude coefficients and handling different levels of sparsity [8]. For the special choice of the free parameters, GBM results in the Horseshoe distribution [9, 10]. The Horseshoe distribution, with a pole at the origin and having heavy tail lead to a sort of nearly optimal choice as a prior for sparse signals. In this work a two-layer hierarchical model is used for representing the GBM and Horseshoe priors. For signal recovery we suggest an iterative algorithm based on the Expectation Maximization method and two fast greedy algorithms inspired by the algorithm in [12]. Our algorithms while enjoy the inherent Bayesian framework advantages, by modeling non-zero coefficients more efficiently and imposing sparsity more heavily, result in lower reconstruction error and faster convergence rate in comparison to the other existing algorithms in this framework. In summary our main contribution are:

- Proposing GBM and Horseshoe as priors in the Bayesian compressive sensing framework.
- Providing an explicit, algebraic EM-update for two-layer model of GBM and Horseshoe, yielding two efficient algorithms.
- Performing an extensive experimental comparison with state-of-the-art methods, demonstrating superiority on a wide parameter range.

The rest of this paper is as follows: in section 2, we study the Bayesian framework for compressive sensing problem, detail hierarchical models and analyze their sparsity-inducing properties. In section 3 we detail Bayesian inference procedure and propose two fast algorithms. We present simulation results and comparison with the other existing algorithms in section 4 and conclude the paper in section 5.

2 Bayesian Framework

In the Bayesian framework, a probability density function $p(\mathbf{w}|\boldsymbol{\gamma})$ is considered on the sparse signal, which reflects our prior knowledge about the nature of the original signal. In this work we consider $p(\mathbf{w}|\boldsymbol{\gamma})$ to be GBM or Horseshoe distribution. These priors features appealing properties compared to Laplace and student's t-distribution. GBM and Horseshoe are much heavier-tailed, modeling large signals efficiently while shrinking noise-like, small signals toward zero.

The measurement vector is also a statistical process where its distribution depends on the distribution of measurement noise. Since we consider the measurement noise as an independent additive zero mean white Gaussian noise with variance equal to β^{-1} , therefore the

conditional pdf $p(\mathbf{y}|\mathbf{w}, \beta)$ is a multiple of normal distributions.

$$p(\mathbf{y}|\mathbf{w}, \beta) = \mathcal{N}(\mathbf{y}|\Phi\mathbf{w}, \beta^{-1}) \quad (4)$$

In the inference procedure for BCS recovery, we need to obtain the joint pdf of all of the problem components $p(\mathbf{y}, \mathbf{w}, \gamma, \beta)$ which can be decomposed as follows:

$$p(\mathbf{y}, \mathbf{w}, \beta, \gamma) = p(\mathbf{y}|\mathbf{w}, \beta)p(\mathbf{w}|\gamma)p(\gamma)p(\beta) \quad (5)$$

In this model β and γ are called *hyper-parameters* and the assigned corresponding distributions $p(\gamma)$ and $p(\beta)$ are called *hyper-priors*.

2.1 Bayesian Hierarchical Model

The optimization problem in (3) corresponds to the *Maximum a Posteriori* (MAP) estimation of \mathbf{w} while Laplace is considered as a prior distribution on the sparse coefficients.

$$P(\mathbf{w}|\lambda) = \frac{\lambda}{2} \exp(-\frac{\lambda}{2}\|\mathbf{w}\|_1) \quad (6)$$

But since in this formulation, Laplace is not the conjugate prior to the conditional pdf $p(\mathbf{y}|\mathbf{w}, \beta)$ a tractable Bayesian inference may not be performed. To address this problem, Bayesian hierarchical models were proposed [4]. In a hierarchical model, the joint pdf $p(\mathbf{w}|\gamma)p(\gamma)$ determines the whole sparsity-inducing properties of the Bayesian model. The conditional prior pdf $p(\mathbf{w}|\gamma)$ is a product of Gaussian pdfs: $p(\mathbf{w}|\gamma) = \prod_{i=1}^N \mathcal{N}(w_i|0, \gamma_i)$ and the hyper-prior $p(\gamma)$ is chosen to be a suitable distribution. For example in the second stage of hierarchical student's t-distribution [3, 4] a Gamma hyper-prior is assigned to the precision variable while in hierarchical model of Laplace a Gamma distribution is considered on the variance of normal [7].

2.2 Generalized Beta Mixture and Horseshoe

In this section we represent the hierarchical GBM and Horseshoe distributions in the conjugate manner. When *Beta2* or *Beta prime* distribution is coupled with normal distribution as a prior for location, it results in a Generalized Beta Mixture.

$$p(w_i|\gamma_i) = \mathcal{N}(w_i|0, \gamma_i) \quad p(\gamma_i|a, b) = \mathcal{B}(\gamma_i|a, b) \quad (7)$$

Where \mathcal{B} denotes the Beta prime distribution and is defined as

$$\mathcal{B}(\xi|a, b) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} \frac{\xi^{(a-1)}}{(\xi+1)^{(a+b)}} \quad \xi > 0$$

and $\Gamma(\cdot)$ shows the Gamma function. The smaller value of a causes sharper peak at zero and consequently concentrates more on mass near zero, while the smaller value of b yields to a distribution which is heavier-tailed [8]. Beta prime distribution can be expressed as a scaled combination of two Gamma distributions in a two layer hierarchical structure which leads to a three layer model for GBM as follows [8]:

$$p(w_i|\gamma_i) = \mathcal{N}(w_i|0, \gamma_i) \quad p(\gamma_i|a, \lambda) = \mathcal{G}(a, \lambda) \quad p(\lambda|b, 1) = \mathcal{G}(b, 1)$$

where $\mathcal{G}(\cdot)$ represents the Gamma distribution. Figure 1 shows different hierarchical models of GBM.

The Horseshoe distribution is a member of GBM family where $a = \frac{1}{2}$, $b = \frac{1}{2}$, and can be represented in every model mentioned for GBM. What is more, Horseshoe is also well

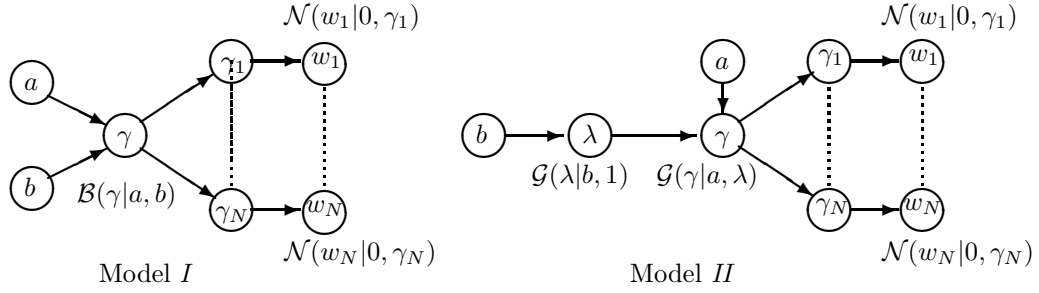


Figure 1: Hierarchical Models of GBM. Model *I* shows the two-layer model and Model *II* shows the three-layer model.

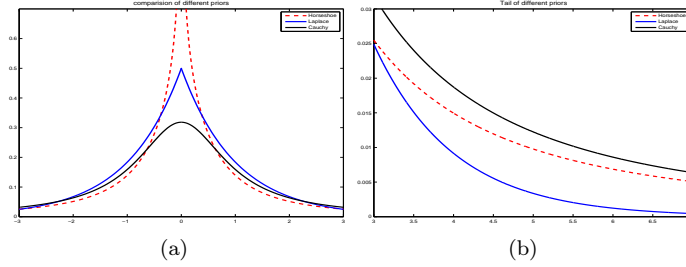


Figure 2: A comparison of Horseshoe versus Cauchy and Laplace distribution. (a) compares different distributions in tail, (b) shows Cauchy, Laplace and Horseshoe distribution at origin.

known as a two-layer hierarchical model in which a half-Cauchy distribution in the second stage is considered on the standard deviation of a normal distribution [8, 9, 10].

$$p(\mathbf{w}|\boldsymbol{\gamma}) = \prod_{i=1}^N \mathcal{N}(w_i|0, \sigma_i^2) \quad p(\sigma) = C^+(0, 1)$$

In our framework we use the two-layer hierarchical model with a Beta prime distribution in the second stage. Horseshoe has an infinite peak near zero which is the result of having a pole at zero. Also it is almost as heavy-tailed as Cauchy distribution [9, 10]. Figure 2 compares the behavior of different distributions in tail and origin.

3 Bayesian Inference

Bayesian inference analysis is based on the posterior distribution. Since the calculation of $p(\mathbf{w}, \boldsymbol{\gamma}, \beta|\mathbf{y})$ is computationally intractable, the exact Bayesian inference is not possible, therefore an approximation method should be utilized. In this work, we drive an inference procedure based on the Expectation Maximization method which iteratively estimates the unknown parameters and signal coefficients. Then by inspiration of the Fast-RVM algorithm [12], we suggest a reconstruction algorithm detailed in Subsection 3.2.2.

3.1 Explicit Algebraic EM-update

We now formulate the Bayesian inference utilizing the EM method. In this regard we decompose the posterior pdf as :

$$p(\mathbf{w}, \boldsymbol{\gamma}, \beta|\mathbf{y}) = p(\mathbf{w}|\mathbf{y}, \boldsymbol{\gamma}, \beta)p(\boldsymbol{\gamma}, \beta|\mathbf{y}) \quad (8)$$

The maximization of the posterior distribution on the left hand side of (8) is almost impossible, hence a two-step approach is adopted. In the first step by maximizing $p(\boldsymbol{\gamma}, \beta|\mathbf{y})$

the MAP estimation of $\{\gamma, \beta\}$ is computed, while the second step calculates the conditional MAP estimation of \mathbf{w} , with $\{\gamma, \beta\}$ set to values obtained in the first step. Here \mathbf{w} is considered as the hidden variable. The E-step of the EM algorithm aims to compute the conditional expectation

$$\mathbf{E}\{\log p(\mathbf{y}, \mathbf{w}, \gamma, \beta)\} \quad (9)$$

with respect to the conditional distribution $p(\mathbf{w}|\mathbf{y}, \gamma, \beta)$ which is a multivariate Gaussian $p(\mathbf{w}|\gamma, \beta, \mathbf{y}) = \mathcal{N}(\mathbf{w}|\boldsymbol{\mu}, \Sigma)$ with parameters [4]:

$$\boldsymbol{\mu} = \Sigma\beta\Phi^T\mathbf{y} \quad \Sigma = (\beta\Phi^T\Phi + \Lambda)^{-1} \quad (10)$$

where, $\Lambda = \text{diag}(1/\gamma_i)$. The M-step updates the estimation of hyper-parameters $\{\gamma, \beta\}$ by calculating the maximizers of (9). During the E-step $\{\gamma, \beta\}$ are set to the value already obtained in the M-step.

By using the relation $p(\gamma, \beta|\mathbf{y}) = \frac{p(\gamma, \beta, \mathbf{y})}{p(\mathbf{y})} \propto p(\mathbf{y}, \gamma, \beta)$ the estimation of $\{\gamma, \beta\}$ is computed by maximizing over $p(\mathbf{y}, \gamma, \beta)$. In order to obtain $p(\mathbf{y}, \gamma, \beta)$, we should marginalize out \mathbf{w} , from the joint distribution as follows:

$$p(\mathbf{y}, \gamma, \beta) = \int p(\mathbf{y}, \mathbf{w}, \gamma, \beta) d\mathbf{w} = \int p(\mathbf{y}|\mathbf{w}, \beta) p(\mathbf{w}|\gamma) p(\gamma|a, b) d\mathbf{w}$$

Consequently the log-likelihood function is :

$$\begin{aligned} \mathcal{L} = & -\frac{1}{2} \log |C| - \frac{1}{2} \mathbf{y}^T C^{-1} \mathbf{y} + (a-1) \sum \log(\gamma_i) \\ & - (a+b) \sum \log(1 + \gamma_i) - N \log(\text{Beta}(a, b)) \end{aligned}$$

Where C is defined as $C = (\beta^{-1}\mathbf{I} + \Phi\Lambda^{-1}\Phi^T)$. We compute the maximizer of log-likelihood function by taking the derivative of above equation with respect to each hyper-parameter and setting the result equal to zero. Doing so we obtain:

$$\begin{aligned} \gamma_i = & \frac{\langle w_i^2 \rangle + 2a - 3}{2b + 3} + \frac{\sqrt{(\langle w_i^2 \rangle + 2a - 3)^2 + 4(2b + 3)\langle w_i^2 \rangle}}{2b + 3} \\ \beta = & \frac{N}{\|\mathbf{y} - \Phi\boldsymbol{\mu}\|^2} \end{aligned}$$

For more details about computing the maximizers of likelihood equation, we refer the reader to section III of [7]. The mentioned EM algorithm suffers from some drawbacks. For example the EM method has low convergence rate and a high dimensional matrix inversion is required at each iteration. To address these problems and decrease the computational requirements a fast suboptimal approach was suggested in [12]. Within this framework and by considering GBM prior for sparse signals, we suggest a fast iterative sequential algorithm which is summarized in Algorithm1 in the next section.

3.2 Fast Greedy Update

The fast algorithm commences with an empty model and sequentially adds the relevant basis. Also as an important feature, in the fast algorithm a basis function can be deleted from the model once it was added [12]. The fast method is relied on efficiently calculating the maximum of the marginal likelihood equation with respect to a single hyper-parameter γ_i . Instead of the $\boldsymbol{\gamma}$ vector, only a single γ_i will be updated at each iteration. In this regard, by following the approach in [12], dependency of $\mathcal{L}(\boldsymbol{\gamma})$ on a single hyper-parameter γ_i should be isolated. So we rewrite the likelihood function as:

$$\begin{aligned}
\mathcal{L}(\gamma) &= [-\frac{1}{2} \log |C^{-i}| - \frac{1}{2} \mathbf{y}^T C_{-i}^{-1} \mathbf{y} + (a-1) \sum_{i \neq j} \log(\gamma_j) - (a+b) \sum_{i \neq j} \log(1 + \gamma_j)] \\
&+ [\frac{1}{2} \log \frac{1}{1 + \gamma_i s_i} + \frac{1}{2} \frac{q_i^2 \gamma_i}{1 + \gamma_i s_i} + (a-1) \log \gamma_i - (a+b) \log(1 + \gamma_i)] + c \\
&= \mathcal{L}(\gamma_{-i}) + l(\gamma_i) + c
\end{aligned}$$

Where C_{-i} is C when the contribution of i^{th} basis is removed and s_i and q_i are defined as :

$$s_i = \phi_i^T C_{-i}^{-1} \phi_i \quad q_i = \phi_i^T C_{-i}^{-1} \mathbf{y}$$

Since C_{-i}^{-1} do not depend on γ_i , s_i and q_i are independent of γ_i . The terms related to γ_i are now isolated. We are looking for γ_i which maximizes the likelihood equation and at the same time we should consider that the definition region for γ_i is \mathbf{R}^+ . The derivative of likelihood function with respect to γ_i is as follows:

$$\frac{d\mathcal{L}(\gamma)}{d\gamma_i} = \frac{dl(\gamma_i)}{d\gamma_i} = -\frac{1}{2} \frac{s_i}{1 + s_i \gamma_i} + \frac{1}{2} \frac{q_i^2}{(1 + \gamma_i s_i)^2} + \frac{a-1}{\gamma_i} - \frac{a+b}{1 + \gamma_i}$$

Equating $\frac{dl(\gamma_i)}{d\gamma_i}$ to zero leads us to a cubic equation with a form of $h(\gamma_i) = m_1 \gamma_i^3 + m_2 \gamma_i^2 + m_3 \gamma_i + m_4$, where

$$\begin{aligned}
m_1 &= -(3 + 2a)s_i^2 \\
m_2 &= (5 + 4b)s_i + (3 - 2a)s_i^2 - q_i^2 \\
m_3 &= (5 - 4a)s_i - q_i^2 + 2 + 2b \\
m_4 &= 2(a - 1)
\end{aligned}$$

The roots of $h(\gamma_i)$ can be determined analytically, but a further analysis is required to realize the maximizer of $l(\gamma_i)$. Note that, depending on the sign of $(a - 1)$, $l(\gamma_i)$ diverges either to plus or minus infinity when γ_i tends to zero. But for $a \geq 1$, GBM is not a sparsity-inducing prior so we limit ourselves to $a < 1$ and continue the inference procedure with this assumption.

3.2.1 The stationary points of $l(\gamma_i)$

Since $l(\gamma_i)$ diverges to plus infinity at the origin, it has no global maximum but we can prove that it has at most one local maximum. In the case of existence, we select the local maximizer of $l(\gamma_i)$, otherwise $\gamma_i = 0$ is the maximizer of $l(\gamma_i)$ and the corresponding basis vector should be pruned from the model.

Proposition 1. *if $a < 1$, $l(\gamma_i)$ has at most one local maximum.*

For $a < 1$ since $h(0) < 0$ and $\lim_{\gamma_i \rightarrow \infty} h(\gamma_i) = -\infty$, $h(\gamma_i)$ has either no positive root or two positive roots. If there exists two positive roots and they are distinct, the greater root is the local maximizer of $l(\gamma_i)$, otherwise the maximum of $l(\gamma_i)$ happens at $\gamma_i = 0$. The above cases can be distinguished through the discriminant of the cubic equation which reveals the nature of the roots. For a cubic equation, the discriminant is defined as $\Delta = 18m_1 m_2 m_3 m_4 - 4m_2^3 m_4 + m_2^2 m_3^2 - 4m_1 m_3^3 - 27m_1^2 m_4^2$. For $\Delta > 0$ the equation has three distinct real roots, for $\Delta = 0$ a multiple root when all roots are real and if $\Delta < 0$ the equation has one real root and two non-real complex conjugate roots. Provided that the discriminant of $h(\gamma_i)$ is positive, it has either three negative distinct roots or two positive and one negative root. The two positive roots happens in two scenarios:

I $\Delta > 0$ and $m_3 > 0$, in this case the stationary points of $h(\gamma_i)$ lies in the different sides of y -axis, so the three roots include positive roots.

II $\Delta > 0, m_3 < 0$ and $m_2 < 0$, in this case the stationary points of $h(\gamma_i)$ have the same sign but since $m_2 < 0$, the sum of them is positive. Hence we conclude that $h(\gamma_i)$ has two positive roots.

3.2.2 Fast Algorithm

Based on the inference procedure detailed in Subsection 3.2, we suggest a greedy algorithm which is summarized in Algorithm 1. The free parameters of GBM, a and b , are set manually at step 1 of the algorithm. In order to reach a higher performance, the free parameters should be tuned by considering the degree of signal sparsity and non-zero coefficients amplitude. At step 5 of the algorithm a basis function should be selected from the basis functions included in or excluded from the model for updating. The selecting procedure could be done randomly or on an order. Following the approach in [12], we calculate the change in marginal likelihood for each potential candidate. The γ_i , which results in the greatest change will be chosen for updating. At each iteration along with γ_i , the value of s_i , q_i , μ and Σ should also be updated. For the sake of computational efficiency, we follow the method in RVM formulation for updating s_i , q_i , μ and Σ in the steps 12 and 13 (see Appendix A of [12] for more details), but due to using new priors the likelihood equation and the updating procedure for γ_i is quite different from the similar algorithms in [7, 12].

Algorithm 1 Fast-GBM

```

1: Inputs :  $\Phi, \mathbf{y}, a < 1, b$ 
2: Outputs :  $\mathbf{w}, \gamma, \Sigma$ 
3: Initiate :  $\gamma = 0$ 
4: while convergence criterion not met do
5:   Choose a  $\gamma_i$ 
6:   if  $\Delta > 0 \ \& \ m_3 > 0 \mid \Delta > 0 \ \& \ m_3 < 0 \ \& \ m_2 > 0$  then
7:     if  $\gamma_i = 0$  then Choose the greater root as the new update for  $\gamma_i$  and add  $\phi_i$  to
       the model
8:     else if  $\gamma_i > 0$  then Choose the greater root as new update for  $\gamma_i$ 
9:     end if
10:  else Prune  $i^{th}$  basis from the model (set  $\gamma_i = 0$ )
11:  end if stop criterion
12:  Upate  $\Sigma, \mu$ 
13:  Upate  $s_i, q_i$ 
14: end while

```

Remark 1. As mentioned in section 2, Horseshoe is GBM distribution with specified hyperparameters ($a = b = \frac{1}{2}$). Since $a = \frac{1}{2}$ and less than 1, the inference procedure for Horseshoe is exactly the same as GBM; therefore we should just simply substitute $a = b = \frac{1}{2}$ in the likelihood equation.

4 Experimental Validation

In this section, we present simulation results to analyze the performance of the proposed algorithms in Subsection 3.2.2. We compare the performance of our algorithms with the other state-of-art sparse estimators in the Bayesian framework, Fast-BCS¹ and Fast-Laplace². We consider a signal of dimension N where T coefficients at random location are drawn from uniform distribution and the other $N - T$ coefficients are set to zero. The measurement matrix

¹<http://people.ee.duke.edu/~lcarin/BCS.html>

²<http://ivpl.eecs.northwestern.edu/>

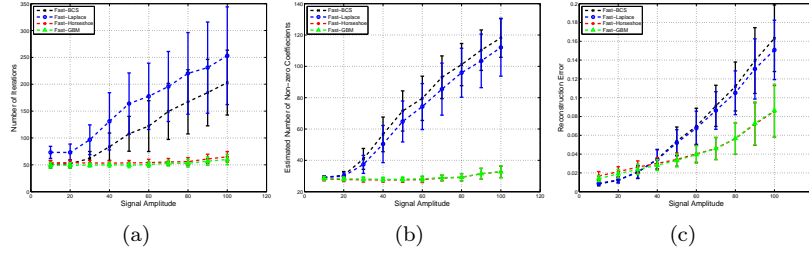


Figure 3: Signal coefficients amplitude (A) versus algorithms performance. (a) shows A versus number of iterations for algorithm convergence, (b) A versus estimated number of non-zero coefficients, (c) A versus mean square reconstruction error.

Φ is chosen to be a uniform spherical ensemble. Through out every set of the experiments, we fix $N = 1000$ and $SNR = 40db$. We also repeat each experiment 100 times and report the average result. As an error measure we take the relative MSE $\|\hat{\mathbf{w}} - \mathbf{w}\|_2^2 / \|\mathbf{w}\|_2^2$ where $\hat{\mathbf{w}}$ and \mathbf{w} are the estimated and original signal vector respectively. We present analyses with respect to three parameters

1. Performance with Respect to Different Levels of Signal Amplitude: In the first set of our experiments, we aim to evaluate the ability of the algorithms in reconstructing signals with different amplitudes. In this regard, we fix $M = 300$ and $T = 30$, while the non-zero coefficients amplitude varies from 0 to A . A is also varied from 10 to 100 in the steps of 10. We report the reconstruction error, estimated number of non-zero coefficients \hat{T} and number of the required iterations for algorithm convergence. The results are demonstrated in Fig.3. As it has been shown in Fig.3(c) for the signals with the amplitude of less than approximately 35, Fast-Laplace and Fast-BCS show better performance in terms of reconstruction error, while for $A > 35$ our proposed algorithms lead to a lower reconstruction error. This is due to the heavy-tailed properties of GBM and Horseshoe which model coefficients with larger amplitudes more efficiently. Fig.3(b) demonstrates that Fast-Laplace and Fast-BCS fail in estimating the number of non-zero coefficients when the amplitude of spikes is larger than 20, while our proposed algorithms give a reasonable estimation of T for all amounts of the signal amplitude in this range. It is obvious from Fig.3(a) that Fast-GBM and Fast-Horseshoe have faster convergence rate.

2. Performance with Respect To Different Number of Measurements: In the second set of our experiments, we study the performance of the algorithms as a function of the number of measurements M . In these investigations, we fix $T = 30$ and vary the number of measurements M from 240 to 350 in steps of 10 and measure the reconstruction error, the estimated number of non-zero coefficients and the required numbers of iterations for algorithm convergence. The results are shown in Fig.4. Here the amplitude of random spikes varies between 0 and 35 which corresponds to where all of the curves meet in Fig.3(c). It means that for this amplitude all of the algorithms show similar performance in terms of reconstruction error. The simulation results reveal that Fast-GBM outperforms the others in terms of reconstruction error except for $M > 330$, that Fast-Laplace and Fast-BCS show better performance. The results in Fig.4(a) demonstrates that the convergence rate of Fast-GBM and Fast-Horseshoe is roughly independent of the number of measurements and clearly our proposed algorithms require fewer iterations. Fig.4(b) shows that the estimated number of non-zero coefficients returned by Fast-GBM and Fast-Horseshoe seem to be almost constant and considerably lower compared to the obtained results by Fast-BCS and Fast-Laplace.

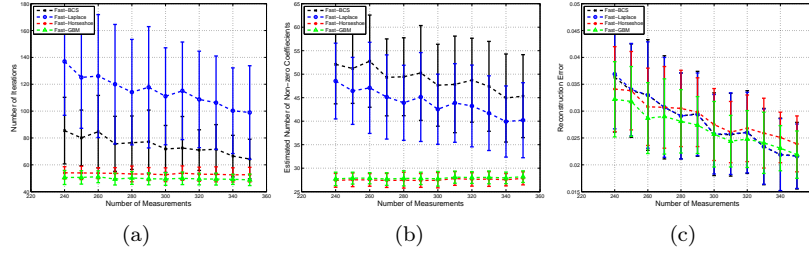


Figure 4: Number of measurements (M) versus algorithms performance. (a) shows M versus number of iterations for algorithm convergence, (b) M versus estimated number of non-zero coefficients, (C) M versus mean square reconstruction error. The amplitude of spikes is between 0 and 35.

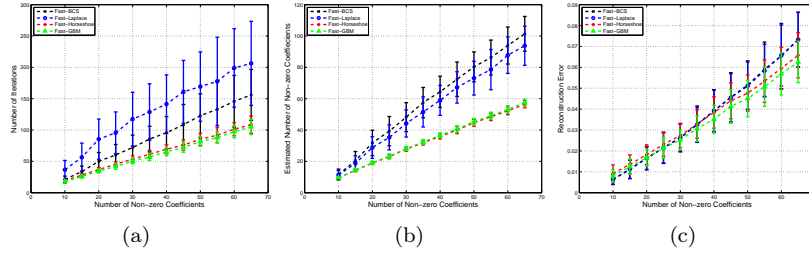


Figure 5: Number of non-zero coefficients (T) versus algorithms performance. (a) shows T versus number of iterations for convergence, (b) T versus estimated number of non-zero coefficients, (C) T versus mean square reconstruction error. The amplitude of spikes is between 0 and 35.

3. Performance with Respect to Different Levels of Signal Sparsity: The third set of experiments is designed to analyze the performance of the algorithms in handling different levels of sparsity. In this regard we vary T from 10 to 65 in the steps of 5, while we fix $M = 300$ and the amplitudes of spikes are between 0 and 35 randomly. The corresponding performance results are depicted in Fig. 5. It can be seen from Fig.5(a) and Fig.5(c) that, while all of the algorithms show similar performance in terms of reconstruction accuracy, our algorithms have faster convergence rate. The results in Fig.5(b) show that Fast-GBM and Fast-Horseshoe considerably outperform the other algorithms in providing an accurate and sparse estimation of T . This is due to the interesting behaviour of the Horseshoe and GBM distributions at origin which enforces sparsity more heavily.

In summary, the simulation results suggest that the performance of our proposed framework is competitive when it is applied to the sparse signals with larger amplitude. However the Fast-GBM and Fast-Horseshoe algorithms provide improved performance in terms of convergence rate and accurate estimation of the number of non-zero coefficients in all of the cases.

5 Conclusion

In this paper we considered the compressive sensing problem in Bayesian framework. We reviewed the existing sparsity-inducing priors in the literature. As new priors, we suggested the Horseshoe and Generalized Beta Mixture as potential candidates for modeling sparse signals in Bayesian structure and studied different hierarchical models for these priors. By considering the GBM and Horseshoe as a prior distribution and using the Expectation Maximization inference method, we provided a solution to CS problem and proposed two reconstruction algorithms resulting from this formulation. The simulation results show our proposed algorithms outperform the other state-of-art algorithms in modeling and estimat-

ing signals with different amplitude and sparsity level. Our algorithms also have faster convergence rate and provide sparser solutions.

References

- [1] Candès, E. J., Tao, T., "Decoding by linear programming," IEEE Trans, Inform. Theory 51 , 4203–4215, 2005.
- [2] J. A. Tropp and A. C. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," Trans. Inf. Theory, vol. 53, no. 12, pp. 4655–4666, 2007.
- [3] S. Ji, Y. Xue, and L. Carin, "Bayesian compressive sensing," IEEE Trans. Signal Process., vol. 56, no. 6, pp. 2346–2356, Jun. 2008.
- [4] M. Tipping, "Sparse Bayesian learning and the relevance vector machine," J. Mach. Learn. Res., pp. 211–244, 2001.
- [5] D. P. Wipf and B. D. Rao, "Sparse Bayesian learning for basis selection," IEEE Trans. Signal Process., vol. 52, no. 8, pp. 2153–2164, Aug. 2004.
- [6] L. He and L. Carin, "Exploiting structure in wavelet-based Bayesian compressive sensing," IEEE Transactions on Signal Processing, vol. 57, no. 9, pp. 3488–3497, 2009.
- [7] S. Babacan, R. Molina, and A. Katsaggelos, "Bayesian compressive sensing using Laplace priors," IEEE Trans. on Image Processing, vol. 19, no. 1, pp. 53–63, 2010.
- [8] Artin Armagan, David B. Dunson, Merlise Clyde, "Generalized Beta Mixtures of Gaussians," Advances in Neural Information Processing System, 2012.
- [9] C. M. Carvalho, N. G. Polson, and J. G. Scott, "The horseshoe estimator for sparse signals," Biometrika, 97(2):465–480, 2010.
- [10] C. M. Carvalho, N. G. Polson, and J. G. Scott, "Handling sparsity via the horseshoe," JMLR, 5, 2009.
- [11] Xiaojing Gu; Leung, H.; Xingsheng Gu, "A variational Bayesian approach to compressive sensing based on Double Lomax priors," Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on , vol., no., pp.5994,5998, 2013.
- [12] M. E. Tipping and A. C. Faul, "Fast marginal likelihood maximization for sparse Bayesian models," in Proc. of the Ninth International Workshop on Artificial Intelligence and Statistics, C. M. Bishop and B. J. Frey, Eds., 2003.
- [13] Carrillo, Rafael E.; Aysal, T.C.; Barner, K.E., "Bayesian compressed sensing using generalized Cauchy priors," Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on , vol., no., pp.4058,4061, March 2010.